COMS 4995: Neural Networks and Deep Learning

Instructor: Richard Zemel

Lectures:	Monday, Wednesday 2:40-3:55
Instructor:	Richard Zemel

Overview

It is very hard to hand design programs to solve many real world problems, e.g. distinguishing images of cats versus dogs. Machine learning algorithms allow computers to learn from example data, and produce a program that does the job. Neural networks are a class of machine learning algorithm originally inspired by the brain, but which have recently have seen a lot of success at practical applications. They're at the heart of production systems at companies like Google and Facebook for image processing, speech-to-text, and language understanding. This course gives an overview of both the foundational ideas and the recent advances in neural net algorithms.

Roughly the first 2/3 of the course focuses on supervised learning – training the network to produce a specified behavior when one has lots of labeled examples of that behavior. The last 1/3 focuses on unsupervised learning and reinforcement learning.

Pre-requisites

This is a second course in machine learning, so it has some substantial prerequisites. Required courses: Machine Learning; Multivariable Calculus; Linear Algebra; Probability & Statistics. These prerequisites will not be enforced, but without them the course will be extremely challenging.

Readings

There is no required textbook for the class. A few small readings may be assigned if the need arises. These required readings will all be available on the web, for free.

There are also some relevant resources which are freely available online. We will try to provide links on a lecture-by-lecture basis.

- Video lectures for UofT Professor Geoffrey Hinton's Coursera course. Professor Hinton is one of the fathers of the field, so think of these as the Feynman Lectures of neural nets.
  `https://www.youtube.com/playlist?list=PLoRl3Ht4JOcdU872GhiYWf6jwrkSNh`
- *Deep Learning*, a textbook by Yoshua Bengio, Ian Goodfellow, and Aaron Courville. `http://www.deeplearningbook.org/`
- Andrej Karpathy's lecture notes on convolutional networks. These are very readable and cover the material in roughly the first half of the course. `http://cs231n.github.io/`
- Richard Socher's lecture notes, focusing on RNNs. `http://cs224d.stanford.edu/syllabus.html`
- Metacademy, an online website which helps you construct personalized learning plans and which has links to lots of resources relevant to particular concepts. We'll post links to relevant Metacademy concepts as the course pro- gresses. `http://www.metacademy.org`
- Video lectures for Hugo Larochelle's neural networks course. These are similar to Professor Hinton's lectures but a bit more mathematical. `http://info.usherbrooke.ca/hlarochelle/neuralnetworks/content.html`
- David MacKay's excellent textbook, *Information Theory, Inference, and Learning Algorithms.* This isn't focused on neural nets per se, but it has some overlap with this course, especially the lectures on Bayesian models. `http://www.inference.phy.cam.ac.uk/mackay/itila/`
- *Neural Networks and Deep Learning*, a book by physicist Michael Nielsen which covers the basics of neural nets and backpropagation.
  `http://neuralnetworksanddeeplearning.com/`

Course requirements and grading

The format of the class will be lecture, with some discussion. I strongly encourage interaction and questions. There are assigned readings for each lecture that are intended to prepare you to participate in the class discussion for that day.

The grading in the class will be divided up as follows:

3 Assignments (Programming & Written)	54%
Mid-Term Test	20%
Project	25%
Attendance & Participation	1%

CLASS SCHEDULE

Shown below are the topics for lectures and tutorials. These are subject to change. The notes will be available on the class web-site the day of the class meeting.

Session	Topic
Lecture 1	Introduction & Linear Models
Tutorial 1	*Multivariable Calculus Review*
Lecture 2	Multilayer Perceptrons & Backpropagation
Tutorial 2	*Autograd and PyTorch*
Lecture 3	Distributed Representations & Optimization
Tutorial 3	*How to Train Neural Networks*
Lecture 4	Convolutional Neural Networks and Image Classification
Tutorial 4	*Convolutional Neural Networks*
Lecture 5	Interpretability
Midterm Quiz	
Lecture 6	Optimization & Generalization
Tutorial 6	*Best Practices of ConvNet Applications*
Lecture 7	Recurrent Neural Networks and Attention
Tutorial 7	*Recurrent Neural Networks*
Lecture 8	Transformers and Autoregressive Models
Tutorial 8	*NLP and Transformers*
Lecture 9	Reversible Models & Generative Adversarial Networks
Tutorial 9	*Information Theory*
Lecture 10	Generative Models & Reinforcement Learning
Tutorial 10	*Generative Adversarial Networks*
Lecture 11	Q-learning & the Game of Go
Tutorial 11	*Policy Gradient and Reinforcement Learning*
Lecture 12	Recent Trends in Deep Learning